

TRANSPORT NETWORK AND METHOD

RELATED APPLICATIONS

This application claims the benefit of U.S. Provisional Application Serial No. 60/202,190, entitled INTERNET PROTOCOL TRANSPORT, filed May 5, 2000 which is
5 hereby incorporated by reference.

TECHNICAL FIELD OF THE INVENTION

The present invention relates generally to the field of telecommunication networks, and more particularly to a
10 transport network and method.

BACKGROUND OF THE INVENTION

Telecommunication networks transport voice and data according to a variety of standards and using a variety of technologies. Circuit-switch networks such as plain
5 old telephone service (POTS) utilize transmission paths dedicated to specific users for the duration of a call and employ continuous, fixed-bandwidth transmission. Packet-switch networks (PSNs) allow dynamic bandwidth, depending on the application, and can be divided into
10 connectionless networks with no dedicated paths and connection-oriented networks with virtual circuits having dedicated bandwidth along a predetermined path. Because packet-switched networks allow traffic from multiple users to share communication links, these networks
15 utilize available bandwidth more efficiently than circuit-switched networks.

Internet protocol (IP) networks are connectionless packet-switched networks. IP networks transport information by breaking up bitstreams into addressable
20 digital packets. Each IP packet includes source and destination addresses and can take any available route between the source and the destination. The IP packets are transmitted independently and then reassembled in the correct sequence at the destination.

25 Traditional IP networks employ synchronous optical network (SONET) or other time division multiplex (TDM) links that support only point-to-point connectivity across the network. Consequently, excessive channelization is required on physical ports of the IP
30 network for point-to-multipoint data flows. This leads to greater cost in terms of support and equipment both in the network and in subtending network elements. In

addition, the number of peering sessions required for subtending routers is significantly increased.

[illegible]

SUMMARY OF THE INVENTION

The present invention provides a transport network and method that substantially eliminate or reduce the problems and disadvantages associated with previous systems and methods. In particular, the transport network provides a flexible topology of ports that are internally configurable into port groups that each act as a single entity and provide point-to-multipoint connectivity between external elements.

In accordance with one embodiment of the present invention, a transport element defined in the transport network includes a port group having a plurality of geographically distributed ports. Point-to-multipoint connectivity is provided between the ports. An identifier is operable to represent the port group as a single element to disparate elements.

More specifically, in accordance with a particular embodiment of the present invention, the identifier is operable to represent the port group as a single element to an external node outside of the transport network. In this and other embodiments, the transport element is operable to participate with the external node in protocol exchanges. The transport element may be a transport router, transport Ethernet switch, transport Frame Relay switch or other suitable element providing point-to-multipoint connectivity between a plurality of external ports.

In another embodiment, routing information for the transport element is generated at a centralized location and distributed to the ports for traffic routing. The routing information may comprise a routing information base (RIB). In this embodiment, the RIB is distributed

to each of the ports, which each generate a forwarding information base (FIB) for the port based on the RIB. To provide a secure data flow in the transport element, the RIB is distributed only to network ports that are part of the port group for the transport element.

In accordance with another aspect of the present invention a node for the transport network includes a first port, a second port, and a processing system. The first and second port each include a receive-transmit pair (RTP) of interfaces. The RTP includes a high-speed demultiplexer operable to process ingress traffic, a high-speed multiplexer operable to process egress traffic, and an external network interface. The processing system is operable to store a first routing model for a first port group including the first port and at least one geographically distributed port. The processing system is also operable to store a second routing model for a second port group including the second port and at least one geographically distributed port. The first port is configurable to process traffic based on the first routing model and the second port is configurable to process traffic based on the second routing model. For security, each port is only allowed access to its own routing model.

In still another aspect of the present invention, a transport element is provisioned by defining at least one port group from a plurality of geographically distributed nodes in the transport network. Point-to-multipoint connectivity is defined between the ports of the port group. The port group is represented as a single entity to an external network for protocol exchanges and traffic processing.

Technical advantages of the present invention include providing an improved transport network. In particular, the transport network provides a flexible topology in which sets of ports may be grouped in any
5 suitable way and treated as a single element. As a result, provisioning of transport elements in the network is unconstrained with mesh and partial-mesh topologies supported. This allows a network owner to freely provision the network to meet customer demands and to
10 optimize efficiency in the network.

Another technical advantage of the present invention includes providing an improved external representation for the transport network. In particular, the transport network is represented as sets of port groups with
15 internally managed connectivity. Port groups are internally defined and each represented as a single entity providing connectivity through the transport network. Accordingly, the external network treats the transport network as a set of distinct elements.

20 Still another technical advantage of the present invention includes providing an improved transport element for the transport network. In particular, the transport element includes a port group with at least point-to-multipoint connectivity between the ports. For
25 mesh topologies, multipoint-to-multipoint connectivity is provided for external point-to-multipoint connectivity. Accordingly, the transport element is externally represented as a router or switch and channelization required on the physical ports is minimized. This
30 reduces cost in terms of support and equipment in the transport network and in subtending network elements.

Still another technical advantage of the present invention includes providing secure transport elements within the transport network. In particular, the topology of the transport network is isolated from the external network. In addition, the routing models and topologies of a transport element is known only to ports of the transport element, and thus hidden from other transport elements in the network. As a result, the transport elements are isolated from one another to prevent cross contamination and intrusions between the elements and customers. Thus, the value of transport elements is increased for both customers and network owners.

Still another technical advantage of the present invention includes providing an improved architecture for nodes in the transport network. In particular, the transport nodes comprise modular building blocks to support a scalable, pay-as-you-grow, architecture. In addition, functionality of the transport elements is distributed between nodes to allow any processor in the transport network to control the element. The control processor for a transport element generates routing information for the element and distributes it to the ports of the element for traffic processing.

Still another technical advantage of the present invention includes providing a transport network and transport elements with high bandwidth efficiency. In particular, optical links are provided between ports of the transport elements in the network. Internet protocol (IP) and other packets are transported directly on the optical links using only a limited or thin transmission protocol. Accordingly, bandwidth use is maximized in the

transport network and the cost of the network is minimized.

Still another technical advantage of the present invention includes providing an improved transport
5 router, transport Ethernet switch, and transport Frame Relay switch with geographically distributed ports. The transport router includes a port group providing point-to-multipoint connectivity between a set of subtending IP routers. The transport Ethernet switch includes a port
10 group providing point-to-multipoint connectivity between a set of subtending Ethernet switches. The transport Frame Relay switch includes a port group providing point-to-multipoint connectivity between a set of subtending Frame Relay switches. In addition, port groups may
15 transparently interconnect any other suitable external ports, in which case the port group may be user protocol independent.

Other technical advantages of the present invention will be readily apparent to one skilled in the art from
20 the following figures, description, and claims.

BRIEF DESCRIPTION OF THE DRAWINGS

For a more complete understanding of the present invention and its advantages, reference is now made to the following description taken in conjunction with the accompanying drawings, wherein like reference numerals represent like parts, in which:

FIGURE 1 is a block diagram illustrating a transport network in accordance with one embodiment of the present invention;

10 FIGURE 2 is a block diagram illustrating an external representation for the transport router of FIGURE 1 in accordance with one embodiment of the present invention;

FIGURE 3 is a block diagram illustrating details of the Internet protocol transport (IPT) node of FIGURE 1 in accordance with one embodiment of the present invention;

15 FIGURE 4 is a block diagram illustrating details of the receiver-transmitter pair (RTP) of FIGURE 3 in accordance with one embodiment of the present invention;

FIGURE 5 is a block diagram illustrating details of the processing system of FIGURE 3 in accordance with one embodiment of the present invention;

20 FIGURE 6 is a block diagram illustrating distribution of functionality between processors in an exemplary network in accordance with one embodiment of the present invention;

FIGURE 7 is a block diagram illustrating details of the transport network layer one (IPTL1) architecture for the processing system of FIGURE 5 in accordance with one embodiment of the present invention;

30 FIGURE 8 is a block diagram illustrating details of the transport element layer two (IPTL2) architecture for

the processing system of FIGURE 5 in accordance with one embodiment of the present invention;

FIGURE 9 is a flow diagram illustrating a method for provisioning an IPT network in accordance with one
5 embodiment of the present invention;

FIGURE 10 is a flow diagram illustrating a method for defining a transport router in an IPT network in accordance with one embodiment of the present invention;

FIGURE 11 is a flow diagram illustrating a method
10 for generating routing tables for a transport router in accordance with one embodiment of the present invention; and

FIGURE 12 is a flow diagram illustrating a method for processing through traffic in a transport router in
15 accordance with one embodiment of the present invention.

DETAILED DESCRIPTION OF THE INVENTION

FIGURE 1 illustrates a transport network 10 in accordance with one embodiment of the present invention. In this embodiment, the transport network 10 is an Internet protocol (IP) network for transporting IP and Multiple Protocol Label Switch (MPLS) packets. The transport network 10 may be any other packet-switched network operable to route, switch, and/or otherwise direct data packets based on network protocol addresses.

The transport network 10 is a private network connecting geographically distributed segments of an external network 12. The external network 12 includes one or more public and/or private networks such as the Internet, an intranet, and other suitable local area networks (LAN), wide area networks (WAN), and nodes. The external network 12 includes label switch and subtending routers 14, Ethernet switches 16, Frame Relay switches 18 and other suitable routers, switches, and nodes operable to generate and/or transport traffic. The transport network 10 communicates with nodes of the external network 12 in the native protocol of the nodes to communicate traffic and control signaling between the networks 10 and 12.

Referring to FIGURE 1, the transport network 10 includes a plurality of Internet protocol transport (IPT) nodes 30 interconnected by communication links 32. The IPT nodes 30 each include a plurality of ports 34 accessible to the external network 12. As used herein, each means every one of at least a subset of the identified items. The communication links 32 are optical fiber or other suitable high-speed links. The high-speed links are operable to transport traffic at a rate of 5

Gb/s or greater. Preferably, the high-speed links 32 transport traffic at rates of 10 Gb/s or above.

As described in more detail below, the high-speed links 32 connect high speed interfaces of the IPT nodes 5 30 to form fast transport segments (FTS) through the transport network 10. Packets transferred via the FTSS incur very small buffering delay in the network as described in co-owned U.S. Patent Application entitled "Method and System for Transporting Traffic in a Packet- 10 Switched Network", filed June 6, 2000. Packets carried through the ports 34 and between FTSS may incur queuing delay comparable to a normal IP switch.

To optimize bandwidth usage within the transport network 10, packets may be transmitted directly on the 15 high-speed optical links 32 without synchronous optical network (SONET) framing and its associated overhead which imposes a penalty of three to five percent depending on the line rate. In one embodiment, a transport label is added to each packet to generate an internal packet that 20 can be directly transmitted on the optical links 32. Details of the transport label are described in co-owned U.S. Patent Application entitled "System and Method for Connectionless/Connection Oriented Signal Transport", filed June 6, 2000. Using the transport label, both 25 connection-oriented and connectionless traffic may be seamlessly transported across the transport network 10. Protection for connection oriented data flows may be provided as described in co-owned U.S. Patent Application entitled "Method and System For Providing A Protection 30 Path For Connection-Oriented Signals In A Telecommunications Network", filed June 6, 2000. Protection for connectionless, packet transport, traffic

flows may be provided as described in co-owned U.S. Patent Application "*Method and System For Providing A Protection Path For Connectionless Signals In A Telecommunications Network*", filed June 6, 2000.

5 To support voice, video, and other real-time or time-sensitive applications, the transport network 10 may provide class of service (CoS) capabilities. In one embodiment, all IP packets are mapped to one of three priority levels as they enter the transport network 10.

10 In this embodiment, guaranteed traffic has reserved bandwidth and is guaranteed to be transported within a defined time delay. Control flow traffic is also reserved and guaranteed, but the network 10 does not guarantee delivery time delay. Best effort traffic does

15 not have reserved bandwidth and delivery is not guaranteed by the network 10. By distinguishing and prioritizing traffic based on its type, including CoS, service level agreement (SLA) and/or other suitable indication of importance or delivery constraints. The

20 transport network 10 is able to deliver time-sensitive traffic within tight time constraints by delaying and/or dropping best effort traffic and other low priority traffic.

In one embodiment, the transport network 10 utilizes

25 a private internal addressing scheme to isolate the network 10 from customers and thus minimize or prevent conflicts with private and/or public networks connected to the transport network 10. This reduces the complexity of network management and preserves the topology of the

30 existing routed network 12. In addition, transport network isolation enables value added services to be provided through the transport network 10.

When an independent addressing scheme is utilized for the transport network 10, egress traffic is converted from the external addressing scheme to the internal addressing scheme at ports 34 using standardized or extended network address translation (NAT). Similarly, egress traffic is converted from the internal addressing scheme back to the external addressing scheme at ports 34 using standard or extended NAT. In addition to the internal addresses, each IPT node 30, port 34 and other component of the transport network 10 visible to the external network 12 includes a globally unique IP address. These addresses are used for external management of the transport network 10.

The transport network 10 provides a flexible topology in which sets of ports 34 may be grouped in any suitable way and each treated as a single entity capable of independently interacting with external nodes. Thus, the transport network 10 is externally represented as sets of port groups 50 with internally managed connectivity. Provisioning of port groups 50 in the transport network 10 is unconstrained with mesh and partial-mesh topologies supported.

The port groups 50 are each a set of ports 34 with similar routing properties. In particular, a port group 50 is a set of ports 34 configured to provide multipoint-to-multipoint or at least point-to-multipoint connectivity between one another which allows point-to-multipoint connectivity between external elements. Accordingly, traffic received by a port group 50 can be routed directly from an ingress port 34 to a plurality of egress ports 34 without channelization in the transport network 10.

Port groups 50 may be provisioned as simple port groups or as composite port groups. In the simple port group configuration, each port 34 only belongs to a single port group 50. Private addresses can be supported
5 inside the simple port group configuration. A composite port group includes ports 34 which have membership in multiple port groups 50. In the composite port group case, private IP addressing is not supported.

The port groups 50 each define a transport element
10 52 with geographically distributed ports 34. Each transport element 52 is assigned a unique global IP address for peering and protocol exchanges within and/or external to the transport network 10. As described in more detail below, the transport elements 52 may
15 implement a distributed architecture in which local processors control each of the ports 34 and a centralized processor controls the network element 52.

In particular embodiments, the transport elements may be transport routers 60 interconnecting sets of
20 subtending IP routers 14, transport Ethernet switches 62 interconnecting sets of subtending Ethernet switches 16, and transport Frame Relay switches 64 interconnecting sets of subtending Frame Relay switches 18. In addition, the transport element 52 may interconnect two ports
25 transparently, in which case the port group 50 is user protocol independent.

FIGURE 2 illustrates details of the transport router 60 in accordance with one embodiment of the present invention. In this embodiment, the transport router 60
30 comprises a simple port group and acts as a single network element within a customer's autonomous network.

Referring to FIGURE 2, the transport router 60 includes geographically distributed ports 34 connected to external routers 14. The external ports 34 form a port group 50 with point-to-multipoint connectivity between the ports 34 as externally represented by the router 80. Accordingly, traffic from any one of the external routers 14 may be routed from an ingress port 34 directly to any number of the other external routers 14 by router 80.

The transport router 60 includes a router identifier to peer with the external routers 14 and participate in reservation and other protocol exchanges. In a particular embodiment, the transport router 60 peers with subtending routers 14 by using interior gateway protocols (IGP) such as OSPF, IS-IS, or RIP. The transport router 60 may peer using an exterior gateway protocol (EGP) or any other suitable protocol.

FIGURE 3 illustrates details of the IPT node 30 in accordance with one embodiment of the present invention. In this embodiment, the IPT node 30 comprises an add/drop multiplexer (ADM) with modular building blocks to support a scalable, pay-as-you-grow architecture. Accordingly, the transport network 10 owner may add functionality and incur cost based on customer demand.

Referring to FIGURE 3, the IPT node 30 includes one or more receiver-transceiver pairs (RTP) 100 and a processing system 102 interconnected by an internal Ethernet connection. As described in more detail below, each RTP 100 includes one or more internal interfaces 104 and one or more external interfaces 106. The internal interfaces are high-speed interfaces between the IPT nodes 30 while the external interfaces 106 are low-speed ports 34 accessible to external nodes. The internal and

local interfaces 104 and 106 may each be implemented as one or more discrete cards.

Within the transport network 10, a set of internal interfaces 104 of the IPT nodes 30 are connected together
5 between ports 34 of a port group 50 to form an FTS between the ports 34 and provide multipoint-to-multipoint and/or point-to-multipoint connectivity. In particular, a multiplexer of an internal interface 104 is connected to a demultiplexer of a next internal interface 104 in
10 the FTS while a demultiplexer of the internal interface 104 is connected to a multiplexer of a previous internal interface 104 in the FTS. The FTSS are directionally-sensitive to preferentially route pass-through traffic over local ingress traffic. In this way, traffic for a
15 transport element 52 is transported between an ingress and egress port on an FTS with minimal delay across the transport network 10.

The processing system 102 includes one or more central processing units (CPUs) 108. The CPUs 108 may
20 each operate the IPT node 30 or a transport element 52. A CPU 108 operating the IPT node 30 includes an operating system and control functionality for the IPT node 30. A CPU 108 operating a transport element 52 includes control functionality for the distributed components of the
25 transport element 52.

FIGURE 4 illustrates details of the RTP 100 in accordance with one embodiment of the present invention. In this embodiment, the internal interface 104 is a high-speed interface that operates at substantially 10 Gb/s.
30 The external interface 106 is a low-speed packet over SONET (POS) interface that operates at 2.5 Gb/s or below.

Referring to FIGURE 4, the internal interface 104 includes an optical receiver 110, a demultiplexer 112, a multiplexer 114, and an optical transmitter 116. The optical receiver is a 10 Gb/s receiver without SONET or package level knowledge. The optical receiver 110 performs the optical to electrical signal conversion. The optical receiver 110 may include an amplifier and may directly interface with a wave division multiplex (WDM) system.

10 The demultiplexer 112 drops local traffic and inter RTP traffic as well as buffers transit traffic. In a particular embodiment, the demultiplexer 112 has a set of 155 Mb/s connections to interface cards of the external interface 106. The demultiplexer 112 may also have 155 Mb/s connections to interface cards of other RTPs 100.

The multiplexer 114 collects local traffic from the interface cards of the external interface 106 and through traffic from the demultiplexer 112. The multiplexer 114 includes packet buffer, scheduler and insertion control 20 functionality.

The optical transmitter 116 is a 10 Gb/s transmitter without SONET or package level knowledge. The optical transmitter 116 may include an optical amplifier. The optical transmitter 116 performs a conversion from an electrical signal to an optical signal and may interface 25 directly with a WDM system.

The external interface 106 include a plurality of low-speed interface cards 120. The low-speed interface cards 120 send and receive traffic to and from the multiplexer 114 and demultiplexer 112, respectively. The low-speed interface cards 120 also provide connections 30 between the FTSS.

The low-speed interface cards 120 are the main buffering point for ingress and egress traffic of the transport network 10. Packet level intelligence, including routing and protection mechanisms, are provided by the low-speed interface cards 120. If the transport network 10 uses an isolated addressing scheme, the low-speed interface cards 120 perform NAT functionality.

FIGURE 5 illustrates details of the processing system 102 in accordance with one embodiment of the present invention. In this embodiment, the transport network 10 includes an internal (IPTL1) layer and an external (IPTL2) layer. The processing system 102 provides a distributed architecture for the transport element 52. In particular, each port 34 of a transport element 52 is locally managed with control processing performed by a centralized processor.

Referring to FIGURE 5, the processing system 102 includes four CPUs 108 each configurable to operate the IPT node 30 or a transport element 52. The first CPU 140 manages the IPT node 30 and includes a simple network management protocol (SNMP) agent/internal network layer one (IPTL1) management information base (MIB) 142 for the IPT node 30. A common management information base (CMIB) 144 includes a model 146 of the transport network 10 and slave models 148 for transport elements having local ports. A database manager 150 manages the CMIB 144. An internal transport network layer one (IPTL1) architecture 152 includes an internal open shortest path first (IOSPF) instance 154 for discovery of the transport network 10. The IPTL1 architecture also includes control component subsystems 156.

The second CPU 160 is a master controller for a first transport element 52 of the transport network 10. The second CPU 160 includes an SNMP agent/external network MIB 162 for the first transport element 52. A
5 CMIB 164 includes a master model 166 of the layer two (IPTL2) architecture for the first transport element 52. A database manager 168 manages the CMIB 166. The IPTL2 architecture 170 includes an OSPF instance 172 for discovery of the network connected to the first transport
10 element 52. The IPTL2 architecture also includes control component subsystems 174.

The third CPU 180 is a master controller for a second transport element 52 of the transport network 10. The third CPU 180 includes an SNMP agent/external network
15 MIB 182 for a second transport element 52. A CMIB 184 includes the master model 186 of the IPTL2 architecture for the second transport element 52. A database manager 188 manages the CMIB 184. The IPTL2 architecture 190 includes an OSPF instance 192 for discovery of the
20 network connected to the second transport element 52. The IPTL2 architecture also includes control component subsystems 194.

The OSPF instances for each transport element discovers the topology for the element and generates the
25 master model. The model is then distributed to the port controllers as slave models for point-to-multipoint connectivity within the port group of the transport element. The fourth CPU 198 is unassigned to a particular transport element 52 and may be idle or used
30 to control lower layer functions.

In operation, layer one (IPTL1) learns the internal topology and does not exchange this information outside

the transport network 10. The internal paths are learned using IPTL1 in order to route traffic between any two points within the network 10 regardless of the contents of the package. The traffic may be locally or externally
5 generated. All IPT nodes 30 participate in IPTL1. Layer two (IPTL2) deals with the external topology for a transport router.

Each IPT node 30 is assigned a unique internal OSPF (IOSPF) router identifier. The transport network 10 runs
10 IOSPF between the IPT nodes 30 to provide normal and protection paths between ingress points of the network. As a result, the transport network is modeled as a collection of routers interconnected by point-to-point links.

15 As described in more detail below, path label calculation (PLC) interacts with the IOSPF in order to learn the transport network 10 topology. Based on the learned topology, PLC determines the normal and protection paths. PLC also addresses overlapping paths.
20 After PLC has learned the transport network topology, PLC signals IPTL2 to start running. When IPTL2 converges, OSPF is updated in the forwarding table for the corresponding transport element 52. PLC then populates the look-up table for the ports 34 of the transport
25 element 52.

FIGURE 6 is a block diagram illustrating the distributed control architecture for transportation routers 60 in an exemplary network. The exemplary network includes a first IPT node 200, a second IPT node
30 202, a third IPT node 204, and a fourth IPT node 206.

The first IPT node 200 includes a first and second port for a first transport router, a first port for a

second transport router, and a fourth and fifth port for
a third transport router. The first CPU 210 includes
control functionality for the first IPT node 200 as well
as slave models of the first, second, and third transport
5 routers for controlling the local ports. The second CPU
212 is a master controller for the first transport
router.

The second IPT node 202 includes a third port of the
third transport router and a third and fourth port of the
10 second transport router. The first CPU 220 includes
control functionality for the second IPT node 202 and
slave models of the second and third transport routers
for controlling the local ports. The second CPU 222 is a
primary controller for the third transport router.

15 The third IPT node 204 includes the fourth port of
the first transport router, a second port of the second
transport router, and a first and second port of the
third transport router. The first CPU 230 comprises
control functionality for the third IPT node 204 and
20 slave models of the first, second, and third transport
routers for managing the local ports. The second CPU 232
includes a master controller for the second transport
router.

The fourth IPT node 206 includes a third port of the
25 first transport router and a fifth port of the second
transport router. The first CPU 240 includes control
functionality for the fourth IPT node 206 and slave
models of the second transport routers for controlling
the local ports. In this way, each IPT node and ports of
30 the IPT node are locally managed. The distributed
transport elements are managed by a centralized
controller on any one of the IPT nodes.

FIGURE 7 illustrates the IPTL1 architecture 250 in accordance with one embodiment of the present invention. FIGURE 8 illustrates the IPTL2 architecture 260 in this embodiment in which the transport network 10 uses a transport label to efficiently transport traffic in the network 10. OSPF uses opaque link state advertisements (OLSAs) in order to discover the external network topology.

Referring to FIGURE 7, the functionality of the PLC 252 is based on whether the processor is managing an instance of IOSPF. An IPT node 30 will have only one instance of IOSPF, but each processor will have an instance of PLC 252. The PLC 252 instance associated with IOSPF builds a local configuration database (LDB) from IPTL1 and IPTL2 provision values, creates the OLSA entry from the configuration of IPTL1, tunnels the OLSA entry to IOSPF, retrieves the OLSA database from IOSPF upon IOSPF's notification of convergence, synchronizes the OLSA database with its PLC peers within an IPT node, signals IPTL2 to start by adding the transport router's port IP address, the multicast host, and transport router's IP address to the port prefix table and adding the CPU's label to the transport table of the port. The PLC 252 also receives the IPTL2 forwarding table (IP forwarding table), populates the prefixes, the transport labels and the destinations mapping tables for the ports of the IPTL2.

The PLC 252 receives fault signal from a fault manager which indicate the link failure identifier. In response to a link failure, the PLC 252 determines which label is effected by the link failure and marks the label as invalid in the transport label's table per port. If

the link identifier is local, the OLSA conveys the failure and hands failure processing over to IOSPF.

The PLC 252 also translates an internal reservation protocol (RSVP) request on a normal path. The internal
5 RSVP specifies the ingress and egress ports. The normal path includes a control path and a data path. A control path is a list of IPT nodes 30 to be traversed from a source to a destination. The data path is a list of high speed and slow speed links to be traversed between the
10 source and the destination. If the internal RSVP succeeds in making a reservation on the normal path, it indicates to the PLC 252 the new QoS of the path. The PLC 252 updates the QoS of the normal transport label for the port 34. The same process occurs for the protection
15 path. If the port 34 is not local to the PLC 252, the PLC 252 tunnels the information to the PLC 252 where the port resides to do the update. Further information regarding the internal reservation process is described in co-owned U.S. Patent Application entitled "System and
20 Method for Opaque Application Object Transport", filed June 6, 2000.

The PLC 252 further supports a proprietary MIB for port lookup table and receives requests from MPLS. The requests include an IP destination prefix and an ingress
25 port. The PLC 252 returns the pointers of the normal and protection transport labels and a next-hop IP address of the subtending router 14. The PLC 252 supports a device driver API to update the forwarding table in the port and supports a label translator to reach any point in the
30 transport network 10.

The PLC 252 instance not associated with IOSPF builds a local configuration database (LDB) from IPTL1

and IPTL2 provisioned values, synchronizes the OLSA database with its IOSPF's PLC peers within an IPT node, signals IPTL2 to start by adding the transport router's port IP address, the multicast host, and transport router IP address to the port prefix table and adding the CPU's label to the transport table of the port, populates the prefixes, the transport labels, and the destinations mapping tables for the ports of the IPTL2.

The PLC 252 also receives fault signal from a fault manager which will indicate the link failure identifier. In this case the PLC 252 determines which label has been effected by the link failure and marks the label as invalid in the transport label's table per port.

The PLC 252 further translates an external IP address to IPTL2 to an egress port for external RSVP, receives signals from a PLC 252 associated with IOSPF to update the local port and receives an internal RSVP request on a normal path. As previously described, the internal RSVP will specify the ingress and egress ports. The normal path includes a control path and a data path. The control path is a list of IPT nodes to be traversed from a source to a destination. The data path is a list of high speed links and low speed links to be traversed between the source and the destination. If the internal RSVP has succeeded in making reservation on the normal path, it indicates to the PLC 252 the new quality of service (QoS) of the path. The PLC 252 updates the QoS of the normal transport label for the port. The same process occurs for protection path. The PLC 252 also supports a device driver API to update forwarding table in ports and supports a label translator to reach any point in an transport network. To perform the

necessary functions, IOSPF will include an API to permit the PLC 252 to pass the OLSA to the IOSPF, signal the PLC to retrieve OLSA database, modify OSPF link state database's structure to store and flood OLSA.

5 Referring to FIGURE 8, the IPTL2 architecture 260 comprises the topology for the transport router 60. The transport router manages the ports 34 in its ports group 50. The subtending routers 14 view the transport router 60 as a single router. The transport router 60 reacts to
10 both external and internal changes in topology, which triggers updates between the subtending routers 14 and the transport router 60. Changes inside the transport network 10 that do not impact the states of the port 34 are not reported to the subtending routers 14.

15 As previously described, a master transport router instance resides in a single processor 262 within the transport network 10. Slave processors 264 resides on each transport node 30 including a port 34 for the transport router 60. Each processor 262 and 264
20 associated with the transport router 60 has a port group communication module 266.

A TCP connection is established between the transport routers instance and the ports instances. This connection is used to traffic control data between the
25 transport router 60 and the subtending routers 14. The communication instance for the transport router 60 monitors the states of the transport routers ports 34 via the TCP connection with the ports instance, downloads a forwarding table upon notification from the routers OSPF,
30 requests from the PLC 252 to translate a port 34 to a transport label, interacts with CMP 268 to send and receive packets, and tunnels the management's control

packets to the transport routers ports 34. The ports communication instance establishes TCP connections with the transport router 60, tunnels all control packets to the transport router 60, request from the PLC 252 to
5 translate a port 34 to a transport label, receives a forwarding table from the transport router 60 and downloads a forwarding table to the PLC 252.

FIGURE 9 illustrates a method for provisioning transport elements 52 in the transport network 10 in
10 accordance with one embodiment of the present invention. The method begins at step 350 in which connections are provisioned between the IPT nodes 30. The connections define the FTSS within the transport network 10. At step 352, addresses for each transport elements 52 are defined
15 within the address space for the IPT network 10.

Proceeding to step 354, the internal topology of the transport network is discovered. At step 356, transport elements 52 are defined within the transport network 10. The transport elements 52 each comprise a port group 50
20 and may be a transport router, transport Ethernet switch, or transport Frame Relay switch. At step 358, topology of the transport elements 52 and connected external nodes are discovered.

Next, at step 360, the transport elements 52 each
25 peer with the subtending routers 14 or other external nodes. At step 362, the transport elements 52 generate routing tables for receiving and transmitting packets to and from the external network and within the transport network 10. In this way, the transport elements 52 are
30 freely defined within the transport network 10 to match the topology of the network 10 to needs of customers.

FIGURE 10 illustrates a method for defining a transport element 52 in the transport network 10 in accordance with one embodiment of the present invention. The method begins at step 400 in which a master, or
5 primary processor for the transport element 52 is assigned within the transport network 10. As previously described, the master processor controls the transport element 52 directly and through slave processors local to each of the ports 34. Next, at step 402, ports 34 are
10 identified and assigned to the transport element 52.

Proceeding to step 404, a local processor is assigned or otherwise provided for each port 34 of the transport element 52. In one embodiment, the local processor by default is a master processor for each
15 corresponding IPT node 30. At step 406, an identifier is assigned to the transport element 52 to allow the transport element 52 to participate in protocol exchanges and otherwise appear as a single element to external nodes.

FIGURE 11 illustrates a method for generating routing tables for a transport element 52 in accordance with one embodiment of the present invention. The method begins at step 450 in which a routing information base (RIB) is generated by a master processor for a transport
25 element 52. The RIB is generated based on the IPTL1 and IPTL2 architectures.

At step 452, the RIB is distributed to each port 34 of the transport element 52. At step 454, a forwarding information base (FIB) is generated at each port 34 based
30 on the RIB. The ports 34 use the RIB to process traffic received from the transport network 10 or the external network 12. Step 454 leads to the end of the process by

which routing information is centrally generated and distributed for the transport element 52.

FIGURE 12 illustrates a method for processing through traffic in a transport element 52 in accordance with one embodiment of the present invention. The method begins at step 500 in which an IP packet is received at an ingress port 34 of a transport element 52. At step 502, a transport label is generated based on the IP address using the FIB for the transport element 52.

Proceeding to step 504, the transport label is added to the IP packet to generate an internal packet. At step 506, the internal packet is transported to an egress port 34 of the transport element 52 on high-speed links based on the transport label.

Next, at step 508, the transport label is removed from the IP packet at the egress port 34. At step 510, the IP packet is transmitted to an external destination element. Step 510 leads to the end of the process by which IP packets are transmitted across the transport network 10 on high speed links using transport labels overhead.

Although the present invention has been described with several embodiments, various changes and modifications may be suggested to one skilled in the art. It is intended that the present invention encompass such changes and modifications as fall within the scope of the appended claims.